

01-31-00



Customer No. 20350

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(303) 571-4000

**ASSISTANT COMMISSIONER FOR PATENTS
BOX PATENT APPLICATION
U. S. Patent and Trademark Office
Washington, D.C. 20231**



Specified herewith for filing under 37 CFR 1.53(b) is the
 X] patent application of
] continuation patent application of
] divisional patent application of
] continuation-in-part patent application of

Inventor(s)/Applicant Identifier: Jerome D. Toporek, a citizen of the USA, residing at 343 Rennie Ave. #8, Venice, CA 90291; David C. Palter, a citizen of the USA, residing at 10933 Wellworth Avenue, #3, Los Angeles, CA 90024; Jeremy A. McCooey, a citizen of the USA, residing at 1002 Tiverton Avenue, Los Angeles, CA 90024; Marc B. Hasson, a citizen of the USA, residing at 1722 Goodman Avenue, Redondo Beach, CA 90278; Timothy W. Hartrick, a citizen of the USA, residing at 11965 Dorothy St., Los Angeles, CA 90049; and Kay A. Guyer, 10933 Wellworth Avenue, #11, Los Angeles, CA 90024.

For: INTERNET OVER SATELLITE APPARATUS

X] This application claims priority from each of the following Application Nos./filing dates:
 U. S. Provisional Patent Application Ser. No. 60/118,227, filed February 2, 1999

the disclosure(s) of which is (are) incorporated by reference.
 X] The following two commonly-owned co-pending applications also are incorporated by reference in their entirety for all purposes: U. S. Pat. App. Ser. No. 09/243,185, filed Feb. 2, 1999 and U. S. Pat. App. Ser. No. 09/243,554, filed Feb. 2, 1999.

Enclosed are:

X] 25 page(s) of specification
 X] 4 page(s) of claims
 X] 1 page of Abstract
 X] 11 sheet(s) of [] formal drawing(s).
 X] An assignment and recordation cover sheet of the invention to Mentat Inc., 1145 Gavley Ave., Suite 315, Los Angeles, CA 90024.
 X] A [] signed [] unsigned Declaration & Power of Attorney
 X] A [] signed [] unsigned Declaration.

(Col. 1)	(Col. 2)	
FOR:	NO. FILED	NO. EXTRA
BASIC FEE		
TOTAL CLAIMS	21 - 20	= *1
INDEP. CLAIMS	4 - 3	= *1
<input checked="" type="checkbox"/> X] MULTIPLE DEPENDENT CLAIM PRESENTED		

* If the difference in Col. 1 is less than 0, enter "0" in Col. 2.

SMALL ENTITY		OTHER THAN SMALL ENTITY	
RATE	FEES	RATE	FEES
	\$345.00		\$690.00
x \$9.00 =		x \$18.00 =	\$ 18.00
x \$39.00 =		x \$78.00 =	\$ 78.00
+ \$130.00 =		+ \$260.00 =	
TOTAL		TOTAL	\$786.00

Please charge Deposit Account No. 20-1430 as follows:

X] Filing fee
 X] Any additional fees associated with this paper or during the pendency of this application.
] The issue fee set in 37 CFR 1.18 at or before mailing of the Notice of Allowance, pursuant to 37 CFR 1.311(b)

] A check for \$ _____ is enclosed.
 2 extra copies of this sheet are enclosed.

Respectfully submitted,
 TOWNSEND and TOWNSEND and CREW LLP

Roger T. Barrett
 Reg. No.: 41,599
 Attorneys for Applicant

Telephone: (303) 571-4000 Facsimile: (303) 571-4000

PATENT APPLICATION
INTERNET OVER SATELLITE APPARATUS

Inventors:

Jerome D. Toporek
343 Rennie Ave #8
Venice, CA 90291; and
a citizen of the UNITED STATES OF AMERICA

David C. Palter
10933 Wellworth Avenue #3
Los Angeles, CA 90024
a citizen of the UNITED STATES OF AMERICA

Jeremy A. McCooey
1002 Tiverton Avenue
Los Angeles, CA 90024
a citizen of the UNITED STATES OF AMERICA

Marc B. Hasson
1722 Goodman Avenue
Redondo Beach, CA 90278
a citizen of the UNITED STATES OF AMERICA

Timothy W. Hartrick
11965 Dorothy St.
Los Angeles, CA 90049
a citizen of the UNITED STATES OF AMERICA

Kay A. Guyer
10933 Wellworth Avenue #11
Los Angeles, CA 90024
a citizen of the UNITED STATES OF AMERICA

Assignee:

Mental Inc.
1145 Gayley Avenue, Suite 315
Los Angeles, CA 90024

Entity: Small Business

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(650) 326-2400

INTERNET OVER SATELLITE APPARATUS

5

CROSS-REFERENCES TO RELATED APPLICATIONS

The subject application claims priority from U.S. Provisional Patent Application No. 60/118,227, filed February 2, 1999, in the name of Jerome D. Toporek et. al, titled, "Internet Over Satellite Apparatus," (Attorney Docket Number 16625-001100), the complete disclosure of which is incorporated herein by reference.

10

The following two commonly-owned co-pending applications also are hereby incorporated by reference in their entirety for all purposes:

1. U.S. Patent Application Serial No. 09/243,185, filed February 2, 1999, in the name of Jerome D. Toporek et. al, titled, " Internet Over Satellite System," (Attorney Docket Number 16625-001200); and

15 2. U.S. Patent Application Serial No. 09/243,554, filed February 2, 1999, in the name of Jerome D. Toporek et al, titled, "Internet Over Satellite Method," (Attorney Docket Number 16625-001300)

20

BACKGROUND OF THE INVENTION

The present invention relates to telecommunications devices and techniques. More particularly, the present invention provides a technique including an apparatus for transmitting information using a TCP connection over a satellite system for use in a wide area network of computers such as the Internet. The present apparatus provides a high speed and quality transmission of signals to a distant location over a satellite system using an Xpress Transport Protocol (herein "XTP") protocol, for example. But it will be recognized that the invention has a much wider range of applicability; it can also be applied to a private bridged network, terrestrial wireless network links, and the like.

25 The Internet is an international "super-network" connecting together millions of individual computer networks and computers. The Internet is generally not a single entity. It is an extremely diffuse and complex system over which no single entity has complete authority or control. Although the Internet is widely known for one of its

ways of presenting information through the World Wide Web (herein "Web"), there are many other services currently available based upon the general Internet protocols and infrastructure.

The Web is generally easy to use for people inexperienced with computers.

- 5 Information on the Web can be presented on "pages" of graphics and text that contain "links" to other pages either within the same set of data files (i.e., Web site) or within data files located on other computer networks. Users often access information on the Web using a "browser" program such as one made by Netscape Communications Corporation of Mountain View, California or Microsoft Corporation of Redmond, Washington.
- 10 Browser programs process information from Web sites and display the information using graphics, text, sound, and animation. Accordingly, the Web has become a popular medium for advertising goods and services directly to consumers.

The Internet supports many other forms of communication. For example, the Internet allows one-to-one communication via electronic mail, which is commonly known as "e-mail." The Internet also has "bulletin boards" which are used by users to post colorful text and graphics for others to read and respond to. For real time communication, the Internet has "chat rooms" and the like, which allow users to communicate to each other by way of typed messages. In some cases, the Internet can also be used for voice communication between users. All of these forms of communication are possible, at least in part, by way of an important connection, which allows information to be transmitted from many different servers on the Internet.

- 25 The Internet is based on the TCP/IP protocol suite. At the network layer, and Internet Protocol, which is known as IP, provides a mechanism to deliver packets addressed to individual computers. The Internet has a plurality of computer systems implementing IP, which are generally interconnected to each other using local area networks and dedicated transmission lines to form a wide area network of computers. IP packets are delivered on a best effort basis and are not guaranteed to arrive at their intended destination, which is known as "unreliable" service.

For many applications that require reliable data delivery service, the

- 30 Internet uses a connection mechanism called Transmission Control Protocol, which is known as TCP. TCP has a variety of desirable characteristics that allows it to be suitable for communication over the Internet. For example, TCP considers the data size or best sized segment of data to send from a transmitting server over the communication medium

to a receiving server. Additionally, TCP maintains a timer, which waits for the receiving server to send an acknowledgement of the reception of the data segment from the transmitting server. If the timer times out before the acknowledgement is received, TCP resends the data segment, which provides reliability in the communication. TCP also

5 maintains a checksum on its header and data, which monitors any modification in the data in transit. If the data arrives with an invalid checksum, TCP discards the data and does not acknowledge receiving the data. If data segments arrive out of order, TCP can also recombine the segments in order at the receiving server. Furthermore, TCP provides flow control, which only allows a finite amount of data to be sent to a buffer on the receiving 10 server. This prevents a fast server computer from overflowing all the buffers on a slower server computer.

Although TCP has been well suited for transmitting Internet data between a plurality of servers, which are coupled to each other by way of conventional telephone systems and lines with similar characteristics, TCP has many limitations. For example,

15 TCP is suitable for “short hops” but often slows down communication over extremely long distances which can have significant latency (e.g., transcontinental hops via satellite). Here, the flow control and acknowledgment features take too long to send and receive, which make communication slow. Additionally, opportunities for gains in efficiency still exist. In particular, TCP has been found to be slow and cumbersome for 20 transmission of Internet data over a satellite network, for example. Additionally, networking using satellites often encounters significant bit error rates, which leads to further slowing of TCP connections. These and other limitations will be more fully described according to the Figs. below.

25 From the above, it is seen that a more efficient way of transporting Internet services over large geographical regions using wireless communication media is highly desirable.

SUMMARY OF THE INVENTION

According to the present invention, a technique for improving TCP

30 performance over a wireless wide area network is provided. In an exemplary embodiment, the present invention provides an apparatus for converting information using a TCP connection to an Xpress Transport Protocol (herein "XTP") connection, for

example, which is more suitable for transmission over a wireless network system such as a satellite system.

In a specific embodiment, the present invention provides a communication apparatus for transmitting information over a satellite link. The apparatus includes a TCP

5 interface coupled to a satellite gateway interface for providing a bi-directional flow of information, which has data and a header, using a TCP connection. As merely an example, the bi-directional flow of information can be from the Internet or internet-like network of computers or any network using TCP/IP protocols. The apparatus also includes a processor for converting the information, using the TCP connection, into a

10 satellite protocol for transmission over a satellite link. The satellite protocol has suitable characteristics for transmission of such information over a satellite link. The protocol can include, among others, XTP, XTP-like protocols, and the like.

In an alternative embodiment, the present invention provides a communication apparatus for establishing a plurality of connections for transmission of

15 information using a TCP connection over a satellite link. The apparatus includes a TCP interface for forming a first communication connection between a first client (e.g., user computer) to a first satellite gateway. The apparatus also includes a satellite gateway interface for forming a second communication connection between the first satellite gateway to a second satellite gateway that is over a satellite link. Information describing

20 characteristics (e.g., source and destination IP addresses and port numbers) of the first connection is transmitted to the second satellite gateway. A third communication connection is formed between the second satellite gateway and the destination server. A combination of these connections forms a transparent connection from the original first client to the destination server.

25 In another aspect of the present invention, techniques for matching a data rate for information flowing through a predetermined point in a network, such as a satellite, to the data rate capabilities of a remote point on the network are provided. In one embodiment, the present invention provides rate control to maintain a suitable or

30 predetermined data rate over the satellite device. The rate control provides a plurality of queues for buffering incoming data. Incoming data are checked against a predetermined length. Data in excess of the predetermined length are stored in the queues. At predetermined intervals, information is taken from one or more queues and transmitted along the outgoing data path. This processing can enable systems according to the

present invention to provide for data rate control across a particular point of interest along the network, such as a satellite link.

In another aspect of the present invention, a method of converting a communications system for providing transport of packetized information in a TCP

5 protocol between a client and a server to a system for providing transport of packetized information in a satellite protocol between the client and the server includes a step of installing a first gateway operatively disposed to intercept connections from the client to the server. The first gateway is further able to convert the packetized information from TCP protocol to satellite protocol. A step of installing a second gateway capable of 10 establishing a connection with the server and further operatively disposed to convert packetized information from satellite protocol to TCP protocol is also included in the method. The first gateway and the second gateway are able to establish a connection over a common telecommunications link. The combination of these elements can provide a method for converting a system for TCP telecommunications to satellite protocol 15 communications.

Numerous benefits are achieved by way of the present invention over conventional techniques. In a specific embodiment, the present invention provides a higher throughput than conventional systems. Additionally, the present invention provides for a way to transparently improve TCP performance over a satellite network.

20 In other embodiments, the present invention provides for a novel protocol which is suitable for long latency, high loss, asymmetric bandwidth conditions, which are typical of satellite communications. The present invention also provides for relatively easy implementation into a conventional satellite system. Depending upon the embodiment, one or more of these benefits can be present. These and other advantages or benefits are
25 described throughout the present specification and are described more particularly below.

These and other embodiments of the present invention are described in more detail in conjunction with the text below and attached Figs.

BRIEF DESCRIPTION OF THE DRAWINGS

30 Fig. 1 is a simplified diagram of a satellite system according to an embodiment of the present invention:

Figs. 2 and 2A are simplified diagrams of system architectures according to embodiments of the present invention;

Figs. 3A-3E are simplified diagrams of process diagrams according to embodiments of the present invention;

5 Figs. 4-6 are simplified diagrams of experimental data according to embodiments of the present invention; and

Fig. 7 is a simplified diagram of a gateway interface according to an embodiment of the present invention.

10 **DESCRIPTION OF THE SPECIFIC EMBODIMENTS**

Before discussing the various embodiments, it may assist the reader to more fully understand the limitations of using TCP over a satellite network. Some of these limitations are described more fully below.

15 **Congestion Avoidance:** In order to avoid the possibility of congestive network meltdown, TCP generally assumes that all data loss is caused by congestion and responds by reducing the transmission rate. Over satellite links, however, TCP often misinterprets the long round trip time and bit errors as congestion and responds inappropriately. Similarly, the TCP "Slow Start" algorithm, which over the terrestrial infrastructure prevents new connections from flooding an already congested network, over satellites forces an excessively long ramp-up for each new connection. While these congestion avoidance mechanisms are important in routed environments, they are often ill-suited to satellite links.

20 **Data Acknowledgments:** The simple, heuristic data acknowledgment scheme used by TCP generally does not adapt well to long latency or highly asymmetric bandwidth conditions. To provide reliable data transmission, the TCP receiver often constantly sends acknowledgments for the data received back the sender. The sender generally does not assume any data is lost or corrupted until a multiple of the round trip time has passed without receiving an acknowledgment. If a packet is lost or corrupted, TCP can retransmit all of the data starting from the first missing

5 packet. The algorithm may not respond well over satellite networks where the round trip time is long and error rates are high. Further, the constant stream of acknowledgments frequently wastes precious back channel bandwidth. If the back channel bandwidth is small, the return of the acknowledgments to the sender can become the system bottleneck in some cases.

10 **Window Size:** TCP utilizes a sliding window mechanism to limit the amount of data in flight. When the window becomes substantially full, the sender stops transmitting data until it receives new acknowledgments from the destination server. Over satellite networks, where acknowledgments are slow to return, the TCP window size generally sets a hard limit on the maximum data throughput rate. The minimum window size often needed to fully utilize an error-free link, known as the "bandwidth-delay product," is 100 Kbytes for a T1 satellite link and 675 Kbytes for a 10 MBps link, for example. Bit errors can increase the required window size further. However, most implementations of TCP/IP are often limited to a maximum window size of 64 KB and many common operating systems use a default window size of only 8 KB, imposing a maximum throughput rate over a satellite link of only 128 Kbps per connection, regardless of the bandwidth of the data pipe.

15

20

According to the present invention, a technique for providing a TCP connection over a wide area wireless network is provided. In an exemplary embodiment, 25 the present invention provides an apparatus for converting information using a TCP connection to an XTP connection, which is more suitable for transmission over a wireless network system such as a satellite system. Details of the present apparatus are described throughout the present specification and more particularly below.

30 Fig. 1 is a simplified diagram of a satellite system 100 according to an embodiment of the present invention. This diagram is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Among other features, the system 100 includes a satellite network system, which has satellite(s) 101. The satellite 101 receives

and transmits information between a plurality of ground stations 107, 108 or satellite dishes. Each satellite ground station includes a satellite dish, and a satellite modem 109A, 109B or other form of receiver such as a VSAT indoor unit, which is coupled, directly or indirectly to a satellite gateway 111A, 111B.

5 Satellite antenna 108 receives and transmits signals 103 through satellite 101. Satellite antenna 108 connects to satellite gateway 111B that couples to a wide area network such as the Internet 129 or an internet, which is coupled to a server 131. The gateway 111B also couples through satellite modem 109B. The server and the Internet communication occur in a common protocol such as TCP/IP, UDP/IP or the like. The
10 satellite gateway will be described in more detail below.

Satellite antenna 107 receives and transmits signals 105 through satellite 101. Satellite antenna 107 connects to satellite gateway 111A that couples to a local area network such as Ethernet, Token Ring, FDDI, or others 121, which is coupled to computer terminals 123 or clients. Here, satellite gateway couples to laptop 117, which is
15 coupled through modem 119. The client, the laptop, and the local area network use a common connection format such as TCP/IP, IPX/SPX, or the like. The satellite gateway will be described in more detail below.

In a specific embodiment, satellite gateway 111A intercepts a TCP connection from a client and converts data 105 to a satellite protocol for transmission
20 over satellite 101. The gateway 111A also couples through satellite modem 109A, which transmits data to satellite 101. The satellite gateway 111B on the opposite side of the satellite link translates the data in the satellite protocol back to TCP for communications with the server 131. In a specific embodiment, the present apparatus provides improved performance over conventional techniques. Additionally, the present methods remain
25 substantially transparent to an end user and is fully compatible with the Internet or an internet infrastructure. In many, if not all aspects, no substantial changes are required to the client or server. In preferred embodiments, applications continue to function without substantial hardware and/or software modifications.

Although the above has been described in terms of a specific networking
30 diagram, it is possible to implement the present satellite gateway in other configurations. For example, the present satellite gateway can be one “hub” or central gateway that serves a plurality of remote gateways. Each of the remote gateways is connected to the central gateway to create an individual satellite link.

Other common network topologies can also be used. Further, in some embodiments, different telecommunications links can be used to carry forward and backward paths of a connection. For example, a satellite link can be used to carry data on the forward path, while telephone line can be used for the backward path. Other types of 5 telecommunications hardware may be substituted for the example hardware described here without departing from the scope of the present invention.

Additionally, the present satellite gateway is generally described as a stand alone unit. It will be recognized, however, that the present gateway can be implemented or integrated into a client machine. For example, the present gateway can be 10 implemented on a personal computer using a satellite card, which can be inserted into a WindowsTM 98 machine. The present gateway can also be implemented on other operating systems such as Windows NT, MacOS, and Linux. Of course, the exact manner of implementation will depend upon the application.

15 Additionally, the present satellite gateway may be integrated into a standalone satellite modem, VSAT hardware, router or other network devices.

I. PROTOCOL DESIGN

In a specific embodiment, the present apparatus includes a novel satellite protocol, which provides improved throughput for satellite networks. The present 20 protocol can be designed to respond efficiently to typical satellite latency, bit errors, and asymmetric bandwidth conditions. The present protocol also can take advantage of optimizations possible on a point-to-point link with fixed bandwidth. For further information regarding satellite protocols, such as XTP, reference may be had to Tim Strayer, "A Brief Introduction to the Xpress Transport Protocol," which is incorporated 25 herein by reference in its entirety for all purposes. These and other features of the present satellite protocol are described in more detail below.

The present protocol utilizes an efficient selective retransmission algorithm for the acknowledgment of data. Because the hop over the satellite is a point-to-point link with no intermediate routing, any gaps in the packet sequence can be 30 assumed to be data loss due to corruption. The receiving satellite gateway requests retransmission immediately upon detection of the missing data from the transmitting satellite gateway.

The present protocol is substantially free from the frequent acknowledgment features of conventional TCP. In some embodiments, the present protocol does not use acknowledgments as the primary means of identifying lost data. In these embodiments, the present protocol needs only infrequent acknowledgments to 5 confirm data arrival and clear buffers. (Conventional TCP sends a constant stream of acknowledgments over the reverse channel.) The present protocol reduces back channel usage by 70% or more, thereby increasing the performance of networks where limited back channel bandwidth is the system bottleneck.

The present protocol offers adequately large window sizes for transmission 10 of data between the satellite gateways. Because the bandwidth-delay product over the satellite between the satellite gateways is much larger than that from the satellite gateway to the end node, the large present protocol window enables the bandwidth-delay product to window size ratio to remain relatively constant. The present gateway becomes a buffer for the network, allowing high throughput independent of the window size of the clients 15 and servers.

The present protocol generally does not use unnecessary congestion avoidance mechanisms for the hop over the satellite between the satellite gateways. The present system, however, continues to use “Slow Start” and all other standard TCP 20 congestion avoidance algorithms for the terrestrial connections between the present gateways and the end nodes. The present protocol also uses a streamlined handshake mechanism to reduce the number of round-trip times required for connection set-up.

The present apparatus can run over IP for compatibility with routed networks. Alternatively, it can run directly over the link layer for substantially improved 25 performance in most instances. Depending upon the embodiment one or more of these features can be available.

The above has generally been described in terms of desirable features of the present satellite protocol. It would be recognized that many other types of features can be available. Additionally, the present invention does not necessarily require each of the above features. Any combination can be used depending upon the application. One 30 of ordinary skill in the art would recognize other variations, modifications, and alternatives.

II. XPRESS TRANSPORT PROTOCOL

In a preferred embodiment, the present invention provides an "Xpress Transport Protocol," which is commonly known as XTP. XTP provides orthogonal protocol functions for separation of implementation from policy, separation of rate and flow control, explicit first-class support for multicast, and data delivery service independence. XTP has a variety of suitable characteristics for transmission of data over a satellite link, which are described more fully below.

In a specific embodiment, the present protocol includes a set of mechanisms whose functionality is orthogonal. The most notable effect of this is that

10 XTP separates communication mechanism (e.g., data-gram, virtual circuit, transaction, etc) from an error control policy employed (fully reliable through uncorrected). Further, flow and rate control as well as error control can be tailored to the communication at hand. If desired, any set of these control procedures can be turned off.

The present protocol also provides for separation of rate and flow control.

15 In general, flow control operates on end-to-end buffer space. Rate control is a producer/consumer concept that considers processor speed and congestion. TCP does not provide rate control, and combats congestion with slow-start and other heuristics. XTP provides mechanisms for shaping rate control and flow control independently.

In a specific embodiment, the present protocol provides explicit multicast

20 support. Here, multicast is not an afterthought in XTP. Each mechanism used for uni-cast communications is available for multicast use as well.

The present protocol also has data delivery service independence. In particular, XTP is a transport protocol, yet with the advent of switched networks rather than routed networks, a traditional network layer service may not be appropriate in every

25 instance. XTP generally requires that the underlying data delivery service provides framing and delivery of packets from one XTP-equipped host to another. This could be raw MAC or IP or AAL5. XTP also employs parametric addressing, allowing packets to be addressed with any one of several standard addressing formats. Other features of XTP include, among others, implicit fast connection setup for virtual circuit paradigm; key-
30 based addressing lookups; message priority and scheduling; support for encapsulated and convergence protocols; selective retransmission and acknowledgment; and fixed-size 64-bit aligned frame design.

XTP defines the mechanisms necessary for delivering user data from one end system to one or more other end systems. Each XTP implementation maintains the state of each of its communications. Well-defined packet structures, containing user data or control information, are exchanged in order to effect the user data transfer. The control

5 information is used to provide the requested layer of correctness and to assist in making the transfer efficient. Assurance of correctness is done via error control algorithms and maintenance of a connection state machine. Flow and rate control algorithms, certain protocol modes, and traffic shaping information are used to provide the requested quality of service as efficiently as possible.

10 The collection of information comprising the XTP state at an end system is called a context. This information represents one instance of an active communication between two (or more) XTP endpoints. A context should be created, or instantiated, before sending or receiving XTP packets. There may be many active contexts at an end system, one for each active conversation or message.

15 Each context manages both an outgoing data stream and an incoming data stream. A data stream is an arbitrary length string of sequenced bytes, where each byte is represented by a sequence number. The aggregate of a pair of active contexts and the data streams between them is called an association.

20 A context at an end system is initially in a quiescent state. A user in need of communication services requests that the context be placed into the listening state. The context now listens for an appropriate FIRST packet. The FIRST packet is the initial packet of an association. It contains explicit addressing information. The user should provide all of the necessary information for XTP to match an incoming FIRST packet with the listening context.

25 At another end system a user requests communication service from XTP. Since this user will initiate the association, the context moves from a quiescent state to an active state directly. The active context constructs a FIRST packet, complete with explicit addressing information from the user. The FIRST packet is sent via the underlying data delivery service.

30 When the FIRST packet is received at the first host's XTP implementation, the address is compared against all listening contexts. If a match is found, the listening context moves to the active state. From this point forward an association is established, and communication can be completely symmetric since there are two data streams, one in

each direction, in an association. Also, no other packet during the lifetime of the association will carry explicit addressing information. Rather, a unique "key" is carried in each packet that maps the packet to the appropriate context.

Once all data from one user has been sent, that data stream from that user's

5 context can be closed. Sentinels in the form of options bits in a packet are exchanged to gracefully close the connection. Other forms of less graceful closings are possible by abbreviating this exchange. When both users are done, and both data streams closed, the contexts move into the inactive state. One of the contexts will send a sentinel that causes the association to dissolve. At this point, both contexts return to the quiescent state.

10 Generally all of XTP's packet types use a common header structure. All of the information necessary to steer the packet's payload to the proper point of processing is carried in the header. Much of how an XTP context operates is controlled by a set of bit flags that are concentrated in one field in the packet header. Fifteen flags are defined, including bit flags to control connection shutdown, bit flags to control the

15 acknowledgment policy, and bit flags that are markers in the data stream. The remaining bit flags control the end-to-end operating modes. Examples include enabling or disabling error control or flow control, or enabling multicast mode.

XTP flow control is based on 64-bit sequence numbers and a 64-bit sliding window. XTP also provides rate control whereby an end system or intermediate system

20 can specify the maximum bandwidth and burst size it will accept on a connection. A Traffic Segment provides a means for specifying the shape of the traffic so that both end systems and intermediate systems can manage their resources and facilitate service quality guarantees.

Error control in XTP incorporates positive and, when appropriate, negative

25 acknowledgment to effect retransmission of missing or damaged data packets. Retransmission may be either go-back-N or selective. The retransmission algorithms are conservative: only data that is shown to be missing via control messages may be transmitted. This avoids spurious and possible congestion-causing retransmissions. The error control algorithm, while basically conservative, can also be aggressive: a technique

30 for a quick-acting error notification is provided.

XTP also specifies techniques for extending error control to a multicast environment. The error control algorithm in multicast is identical to the uni-cast algorithm, although additional sophistication is required to manage state variables and

achieve continuous streaming. Accordingly, XTP is a preferred satellite protocol according to the present invention.

Although the above has generally been described in using an XTP protocol, it will be understood that other types of protocols can be used. For example, the protocol can be a modified TCP, a modified XTP, and others. Additionally, the protocol can be any that is suitable for transmission over a satellite link. Additionally, one of ordinary skill in the art would recognize other variations, modifications, and alternatives.

Fig. 2 is a simplified diagram of system architecture 200 according to an embodiment of the present invention. This diagram is merely an illustration and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. The system architecture 200 includes at least a client 201, which is coupled to a satellite gateway 203, which transmits and receives information via satellite link 209 from a satellite gateway 205. Satellite gateway 205 couples to server 207. Satellite gateway and server couple to each other through the Internet or an internet-like network of computers.

Client 201 can be represented in multiple layers, including application and physical layers. The layers may include a browser 211. The browser 211 allows a user to communicate information from an application layer to a physical layer for transmission to a gateway. The browser 211 is generally in the application layer, which provides the information. For example, the browser can be one of suitable design made by a company called Netscape Communications Corporation of Mountain View, California or Microsoft Corporation of Redmond, Washington or others. In addition to browsers, other TCP applications, including FTP file transfers, may also be used to communicate between clients and servers.

The information goes through the transport layer (e.g., TCP) 213 and then through the IP layer 215, which is the networking layer. The transport layer provides a flow of data between the client and the gateway. The IP layer handles movement of data comprising packets between the client and the gateway, or other network. The information is sent through the physical layer, which includes a driver 217. The driver 217 transmits the information to the gateway 203 through hardware 219 connected to gateway 203 via a telecommunications link 221, which can be a wire, cable, fiber optic, or other physical communication medium. Alternatively, the driver 217 receives information from the gateway 203 via link 221 and hardware 219. The driver can be a

network operating system with a network interface card in the client computer, for example.

From the physical connection 221, the information traverses to the gateway 203. The gateway has a physical layer 223, which interacts with driver 225.

- 5 The gateway also includes a networking layer 227 and a transport layer 229, which is coupled to a translation module 231. The networking layer 227 determines whether the information can be routed over the satellite protocol. If so, the data is passed up to the transport layer 229. If not, the data is immediately forwarded out to the rate control module 234 for delivery to the satellite link 239 in non-translated form. The translation module 231 converts the information into a satellite protocol 233, which is suitable for use over a satellite link. A rate control module 234 determines whether the information can be passed immediately to the satellite connection or be queued for later delivery. The satellite connection is coupled to a physical layer 237 by a driver 235, which transmits information to and from satellite 209. The information traverses to and from the satellite in a wireless medium, 239, 241.
- 10
- 15

Information is received by the satellite gateway 205, which includes multiple layers, physical and others. The physical layer 243 couples to driver 245. The networking and transport layer include a satellite protocol 247 and a rate control 244.

- The network and transport layers connect to a session layer which comprises translation module 249. The translation module converts the information using the satellite protocol back to TCP/IP. The information traverses through the transport layer (i.e., TCP) 251 and the networking layer (i.e., IP) 253. Information from the satellite gateway 205 enters a physical layer 257 through driver 255. The driver transmits the information to a server 207 over telecommunications link 259. Driver 255 can also receive information from server 207 in the backward path.
- 20
- 25

From the telecommunications link 259, the information is sent to driver 263 in the server 207 via physical layer 261. The information traverses from the driver to a networking layer, which includes IP 265. The information also traverses from the networking layer to a transport layer, which includes TCP 267. The information is delivered to a Web server application 269, for example. Between the server 207 and the satellite gateway 205 is the Internet or another IP network, depending upon the application.

In a specific embodiment, information can flow through a gateway, such as gateway 203, 205 at the network layer, bypassing the transport and application layers altogether. For example, in gateway 203, information can enter from client 201 via telecommunications link 221. Physical layer 223 passes the information to driver 225. In 5 turn, driver 225 passes the information to network layer 227. Network layer 227 can route the information to its destination using IP. The information then flows from network layer 227 to driver 235. Driver 235 interacts with physical layer 237 to pass the information along to satellite 209 via telecommunications link 239.

In a specific embodiment, the translation module converts information 10 using a TCP connection to a suitable connection for transmission over a satellite system. The suitable connection for the satellite is generally resilient to transmission over a wireless network such as a satellite network, e.g., Orion Worldcast™ or PanAmSat™ SpotBytes™ services, and the like. The connection should also be suitable for long latency, high loss, and asymmetric bandwidth conditions. The connection should also be 15 transparent to the end user at the client or server location. The long latency is typically about 200 milliseconds to about 700 milliseconds, per satellite hop but is not limited to such values.

In a specific embodiment, the present translation module converts information using a TCP connection into an XTP, modified TCP or XTP-like protocol for 20 transmission over a satellite link, which is illustrated by Fig. 2A, for example. Here, the TCP module receives information 350 including the TCP connection. The information includes data 333, a TCP header 335, and an IP header 337. The TCP module removes the TCP header from the information such that the information 351 includes substantially all data 333 and passes the data, along with certain TCP header information, to the 25 translation module. The translation module passes the data or portions of data to the satellite protocol module where a satellite protocol header 339 is prepended onto the data. The header is an XTP header or like. The information 352 now includes the XTP header, which is suitable for transmission over the satellite link. Depending on the specific implementation, the XTP header and data may also be prepended with an IP header for 30 transmission over the satellite link. In particular, the information is transferred to the physical layer and the driver. The driver and physical layer send the information to a receiving satellite gateway, which may prepend additional headers.

From the satellite link, the information including the XTP header enters a receiving XTP module. The receiving XTP or satellite protocol module removes the XTP header from the information, leaving data 353. The receiving satellite protocol module passes the data to the translation module, which can be a routing module used to route data to the proper protocol. The translation module then passes the data to the TCP module. The receiving TCP module attaches a TCP header 343 and IP header 341 onto the data to form information that is now ready for transmission over a network to a receiving server destination. The information using the TCP connection traverses through a network to the destination server.

10 In other embodiments, the translation module can also convert a portion of the information including the data and TCP/IP headers to information using the XTP protocol. Alternatively, the translation module can convert more than one segment of information including multiple TCP/IP headers into a single piece of information including the XTP header for transmission over a satellite link. Depending upon the 15 application, the translation module can convert the information including the TCP connection into any one or any combination of the above. Further details of these techniques, and others, are described in more detail below.

Figs. 3A and 3B are simplified diagrams of process diagrams according to embodiments of the present invention. These diagrams are merely illustrations and should 20 not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. In a specific embodiment, the present invention uses a connection process 300, which is illustrated by Fig. 3A. The connection process uses a plurality of separate connections using a "handshaking routine" in the satellite system to provide a transparent TCP connection to an end user. The satellite 25 system can be similar to the one described herein, but is not limited. The present process begins at start, step 301. In a specific embodiment, a plurality of connections are provided. In particular, the present process provides a first connection, step 303. The first connection is a TCP connection between a client and a client side gateway, which interfaces to a satellite link. The present process provides a second connection, step 305, 30 which provides a connection between the client side gateway and a destination gateway over a wireless (e.g., satellite) link. The connection between the wireless link can be any suitable connection such as XTP or XTP-like connection. The present process provides a third connection, step 307, which forms a TCP connection between the destination

gateway to the destination server, which can traverse through the Internet or an internet. Once these connections have been made they are maintained, step 309. The present process ends at step 311, which terminates all three connections.

In a specific embodiment, the present process occurs by way of a selected sequence of steps. Here, TCP SYN packets are intercepted transparently by a satellite gateway. The satellite gateway establishes a new XTP connection across the satellite link to the other satellite gateway. Information, including IP addresses and port numbers, about the requesting client and the destination server is transferred across the satellite link to the other satellite gateway. The destination gateway then sets up a TCP connection with the destination server, using the client's addressing information so that the server sees the TCP connection as being the same as if the client itself had established the connection. The satellite gateway makes routing decisions based upon source and destination addresses in the IP header of packets entering the gateway in order to determine whether the packets should be routed over the satellite telecommunications link. Once the TCP connection to the server is established, the destination gateway communicates back to the sending gateway, which then acknowledges the connection with the client.

In one or more embodiments, this sequence of steps has advantages. In particular, the client does not believe the connection is established until the server has accepted this connection, which tends to reduce or even eliminate false indications of successful connection establishment. Additionally, both the client and server see the connection as happening immediately between the two of them, without detecting that the satellite connection is happening in between the client and the server. That is, the present invention substantially preserves "end-to-end semantics" of the connection. In other embodiments, the connectivity is symmetric. Here, "clients" may exist on either side of the satellite link, and "servers" may exist on either side. Systems on either side may request connections to systems on the other side, and the satellite gateways will intercept the connections.

Fig. 3B illustrates a simplified flowchart of a novel routing process 320 in a specific embodiment according to the present invention that begins with a starting step 321. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. In a step 323, a client host sends a TCP packet destined for a particular

IP address to a satellite gateway. In an example, the satellite gateway can be similar to the one 111A described above, but can also be others. In a decisional step 325, the satellite gateway receives the packet and determines whether the packet is to be carried over an XTP connection prior to transmission. An example of an XTP protocol has been

5 described herein, but should not limit the scope of the claims. Any suitable satellite protocol can be used according to certain embodiments of the present invention. If the destination address of the TCP packet is configured to be converted to the XTP protocol format, then the packet is converted into the XTP protocol in a step 331. Otherwise, the packet is transmitted in its TCP format via an alternative branch to step 327. By way of
10 the present decisional process, the present invention is useful for sites where a single satellite gateway communicates with multiple remote sites, some of which have gateways that are compatible with XTP protocol while other sites do not contain corresponding gateways. Additionally, the present invention allows specific addresses to be configured for protocol translation and other addresses to be configured for no translation, based on
15 policy decisions of the network administrator.

In a specific embodiment, the present invention also provides a rate control step (step 327). In a step 327, processing is performed to maintain a suitable or predetermined data rate over the satellite device. The processing of step 327 may require buffering of packets in order to avoid overloading the satellite link, such as 105 and 103

20 in Fig. 1. This can be accomplished using the process depicted by Figs. 3C-3D, but is not limited to this process. Then, in a step 329, the packet is transmitted along a connection transmitted from the sending satellite gateway across the satellite link. Then, in a step 333, the packet is converted back to TCP, if it was converted to XTP protocol in step 331. Then, in a step 335, the packet is routed to its remote destination. The present invention
25 is especially useful for sites where a single satellite gateway communicates with multiple remote gateways, some of which are compatible with XTP protocol while others of the remote gateways are not. Some configurations may have multiple remote sites per satellite gateway. Some of these remote sites may have a satellite gateway installed, while others do not. Thus, there may be no satellite gateway on the remote side, just a IP
30 routing configuration. The process completes routing (step 335) when the information enters a destination server 131, which transfers the information to a destination client. The process stops (step 337) when the connection is terminated.

In a particular embodiment according to the present invention, TCP connections at times pass a piece of information in the protocol header which must be delivered as soon as possible to the destination side of the connection. This piece of information is known as the "urgent pointer". In select embodiments, the TCP 5 implementation that is part of the client-side satellite gateway, such as for example satellite gateway 203 of Fig. 2, can extract this urgent pointer from the TCP header. However, this operation is not required by the system of Fig. 2 and is not intended to limit the scope of the claims. A TCP module, such as TCP module 229 of Fig. 2, for example, can extract the urgent pointer and pass it to a satellite gateway translation module, which 10 can be satellite gateway translation module 231 of Fig. 2, or an equivalent. The urgent pointer is passed down to a satellite protocol module, such as satellite protocol module 233 of Fig. 2, or an equivalent. The urgent pointer can then be carried in the satellite protocol header to a satellite gateway, such as satellite gateway 205 of Fig. 2, or any other receiving satellite gateway. At the receiving satellite gateway, the urgent pointer can be 15 extracted by a satellite protocol module, such as satellite protocol module 247 of Fig. 2, and passed up to a translation module, such as translation module 249, or another satellite translation module, which can deliver it to a TCP module, such as TCP module 251. The TCP module then incorporates the urgent pointer in its target field in the TCP header for immediate delivery to an end server, such as end server 207. This header for delivery to 20 the server refers to the point in the data stream that has not yet arrived at the second satellite gateway machine. In some embodiments, appropriate urgent pointer processing can alleviate malfunctions.

Fig. 3C illustrates a simplified flowchart of a data rate control process such as data rate control step 327 in Fig. 3B in a specific embodiment according to the present 25 invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3C shows a first decisional step 341 of determining whether queues are empty when an incoming packet arrives. In a presently preferable embodiment, there are two queues, one for queuing high priority (HIPRI) traffic and one 30 for queuing "normal" traffic. A person of ordinary skill in the art can appreciate that the number of queues can be extended or reduced by straightforward extensions of the embodiments according to the present invention. If the queues are empty, then

processing continues with a decisional step 343. Otherwise, processing continues with a decisional step 345.

Decisional step 343 determines whether the current system clock tick count is the same as a stored clock tick count. If the system clock tick and the stored

- 5 clock tick are the same, then processing continues with decisional step 349, which determines whether a byte counter and a length of an incoming packet are greater than a predetermined maximum, based on the desired transmission rate, denoted here as “MAX”. If the maximum would not be exceeded, then processing continues with step 355, in which the packet is sent and the byte counter is incremented by the packet length.
- 10 Otherwise, if the maximum would be exceeded in step 349, then processing continues with a step 357, in which the difference between MAX and the current value of the byte counter is stored as “EXTRA”, and a timer can be started. Processing then continues at step 345. Otherwise, if step 343 determines that the clock tick and stored clock tick are different, processing continues with a step 347 in which a byte counter is cleared and the
- 15 stored clock tick is updated. Processing then continues with a step 355, as described above.

If decisional step 341 determines that at least one of the queues is not empty, or if the packet length would have increased the byte counter beyond MAX in step 349, then in a decisional step 345, a determination is made whether the incoming packet

- 20 is a retransmission. If the packet is a retransmission, then in a step 351, the packet is added to the high priority (HIPRI) queue. Otherwise, in a step 353, the packet is added to the normal queue.

Fig. 3D illustrates a simplified flowchart of a timer service process useful in conjunction with data rate control process illustrated by Fig. 3C in a specific

- 25 embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 3D shows a step 361 invoked whenever a timer set in step 357 expires. In step 341, the byte counter is cleared and the local maximum is set to the maximum MAX added toEXTRA. Then, in a
- 30 decisional step 363, the high priority (HIPRI) queue is checked to determine whether packets are on the queue. If there are packets on the HIPRI queue, then processing continues with a decisional step 365. Otherwise, processing continues with a decisional step 367.

If there are packets on the HIPRI queue, then in a decisional step 365, a determination is made whether the byte counter and the packet length are less than a local maximum. If this is indeed the case, then processing continues with a step 369 in which the packet is dequeued from the HIPRI queue and sent. Also, the byte counter is
5 incremented to reflect the length of the data in the packet. Otherwise, if decisional step 365 determines that the byte counter and the packet length exceed the local maximum, then in a step 371 the timer is restarted and processing returns to the invoking process. Processing can continue when the timer expires.

If decisional step 363 determines that there are no packets on the HIPRI
10 queue, then processing continues with a decisional step 367. Decisional step 367 determines whether there are packets on the normal queue. If there are no packets on the normal queue, then processing continues with a step 375, in which the current clock tick is stored in a global cell. Otherwise, if decisional step 367 determines that packets are on the normal queue, then processing continues with a decisional step 373, in which a
15 determination is made whether the byte counter and the packet length are less than the local maximum. If this is indeed the case, then processing continues with a step 377, in which the packet is dequeued from the normal queue and sent. Also, the byte counter is incremented to reflect the length of the data in the packet. Otherwise, if decisional step 373 determines that if the byte counter and the packet length exceed the local maximum,
20 then in a step 379 the timer is restarted and processing returns to the invoking process. Processing can continue when the timer expires.

Fig. 3E illustrates a representative system overview in a particular embodiment according to the present invention. This diagram is merely an example which should not limit the scope of the claims herein. One of ordinary skill in the art
25 would recognize other variations, modifications, and alternatives. A client 380 initiates a connection request to a TCP server 388. Satellite gateway 384 intercepts the connection request and establishes a second connection 385 with a second Satellite gateway 386. The second satellite gateway 386 then initiates a third connection 387 with the server 388. Once connection 387 is established and that information has been communicated back to
30 the satellite gateway 384, the first TCP connection 383 can be confirmed between the client 380 and satellite gateway 384. Although depicted in terms of satellite and satellite gateways, the present invention can also be embodied in other forms of network

hardware. For example, first gateway 384 and second gateway 386 can be connected by a wireless network and the like.

Fig. 7 illustrates a representative apparatus in a particular embodiment according to the present invention. This diagram is merely an example which should not

5 limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Fig. 7 depicts a block diagram of a protocol gateway in a particular embodiment according to the present invention, such as protocol gateway 111A, or the like. Protocol gateway 111A includes a bus 112 which interconnects major subsystems such as a central processor 114, a system memory 116 (typically RAM), and a number of optional external devices such as a display screen 124 via a display adapter 126, a keyboard 132 and a mouse 146 via an I/O controller 118, a floppy disk drive 136 operative to receive a floppy disk 138. Storage Interface 134 may act as a storage interface to a fixed disk drive 144. An optional CD-ROM player 140 operative to receive a CD-ROM 142 can also be included. Other forms of storage, such 10 15 20 25 30 as a flash memory 147 can also be included. A network interface 148A can provide a direct connection to a remote server via a telephone link or to the Internet using a common protocol such as TCP/IP or the like. Network interface 148A can interface to one or more outside networks, which can employ any network format known in the art, such as Ethernet, Token Ring, ATM, IEEE 802.3, X.25, Serial Link Internet Protocol (SLIP), FDDI and the like. Network interface 148B may also connect to a satellite gateway, such as satellite gateway 109A, or other network interconnecting devices or computer systems. Network interface 148A can interface to one or more outside networks, which can employ any network format known in the art, such as Ethernet, Token Ring, ATM, IEEE 802.3, X.25, Serial Link Internet Protocol (SLIP), FDDI and the like. Many other devices or subsystems (not shown) may be connected in a similar manner. Also, it is not necessary for all of the devices shown in Fig. 7 to be present to practice the present invention, as discussed below. The devices and subsystems may be interconnected in different ways from that shown in Fig. 7 in various embodiments. Code to implement the present invention, may be operably disposed or stored in computer-readable storage media such as system memory 116, fixed disk 144, CD-ROM 140, or floppy disk 138.

System 111A is merely one example of a configuration that embodies the present invention. It will be readily apparent to one of ordinary skill in the art that many

system types, configurations, and combinations of the above devices are suitable for use in light of the present disclosure.

Although the above has generally described the present invention according to specific systems and embodiments, the present invention has a much broader range of applicability. In particular, the present invention is not limited to satellite telecommunications, but can be applied to any networking situation where an improved or optimized protocol is desired for use over a specific portion of the network, and the end systems cannot be updated to use the improved protocol. Thus, in some embodiments, satellite gateways could provide access to wireless or cabled networks and internetworks of all kinds. Of course, one of ordinary skill in the art would recognize other variations, modifications, and alternatives.

Experiment:

To prove the principles and operation of the present system, experiments have been performed. In the present experiments, we ran a series of single-client FTP file transfer throughput tests to compare performance of conventional TCP and the present invention over a wide variety of different TCP windows sizes, link bandwidths, round trip delay times, and bit error rates. In one experiment, "Window Size" was compared against "Throughput," as shown by the simplified diagram 400 of Fig. 4, for example. This diagram is merely an example and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Without performance enhancements or TCP window size tuning, most clients were limited to a throughput rate of less than 100 Kbps when the round trip time was 540 ms. As the test data in Fig. 4 show, even clients with a 32 KB window were able to reach a throughput of a mere 400 Kbps. The present system allowed the client to take advantage of the available bandwidth regardless of the window size of the client or server.

In an alternative experiment, a "Round Trip Delay" was analyzed against "Throughput," as shown in a simplified diagram 500 of Fig. 5, for example. This diagram is merely an example and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Here, the present invention removed the dependency of TCP on the round trip time of the network. Performance was measured over symmetric 2 Mbps and 10 Mbps links with no errors. These results illustrated that TCP is able to fully saturate terrestrial low-delay

networks, but as the delay increases, TCP performance dropped rapidly. In contrast, a network which used the present system was able to maintain near complete usage of the link regardless of the round trip time of the network.

In an alternative experiment, a "Bit Error Rate" was monitored against

5 "Throughput," which is shown by the simplified diagram 600 of Fig. 6, for example. This diagram is merely an example and should not limit the scope of the claims herein. One of ordinary skill in the art would recognize other variations, modifications, and alternatives. Networks which incorporated the present invention were also sensitive to the bit error rate of the link. The diagram shows throughput as a function of the bit error rate for a
10 symmetric 10 Mbps satellite network with a 540 ms round trip time and a TCP window size of 1 MB. Even at low error rates, TCP was able to deliver a mere 1.2 Mbps. At an error rate of 1.0×10^{-5} , TCP's throughput dropped to 0.05 Mbps. A network which used the present system was able to fully saturate the link at low error rates and delivered 2.7 Mbps even at an error rate of 1.0×10^{-5} .

15 While the above is a full description of the specific embodiments, various modifications, alternative constructions and equivalents may be used. For example, the above has generally been described in terms of using XTP as a satellite protocol. Other types of protocols may be available. For example, the protocol can include a modified TCP or the like. Therefore, the above description and illustrations should not be taken as limiting the scope of the present invention which is defined by the appended claims.
20

WHAT IS CLAIMED IS:

1 1. A communication apparatus for transmitting packetized
2 information, said information comprising a plurality of packets, each of said packets
3 comprising data and a header, over a satellite link in a telecommunications system, said
4 system comprising a client, selected from a plurality of potential clients, a server, selected
5 from a plurality of potential servers, a first gateway, connected to said client by a first
6 telecommunications link, a second gateway, connected to said server by a second
7 telecommunications link, a third telecommunications link connecting said first gateway to
8 said second gateway, said apparatus comprising:

9 a network interface for linking said first gateway with said client;

10 a satellite gateway interface;

11 a system memory; and

12 a bus interconnecting said network interface, said satellite gateway

13 interface, and said system memory with a processor, said processor operatively disposed
14 to:

15 intercept a connection with said server, said communication initiated by
16 said client;

17 establish a connection between said first gateway and said second gateway
18 over said telecommunications link;

19 provide a bi-directional flow of information from said client to said server
20 and from said server to said client using said connection between said first gateway and
21 said second gateway, wherein said providing a bi-directional flow occurs transparently to
22 said client and said server.

1 2. The apparatus of claim 1 wherein said processor is further
2 operatively disposed to convert said information at said first gateway from a first protocol
3 into a second protocol for transmission over said telecommunications link; and

4 convert said second protocol into said first protocol at said second
5 gateway.

1 3. The apparatus of claim 2 wherein the first protocol comprises TCP

2 and said second protocol comprises XTP.

1 4. The apparatus of claim 2 wherein said second protocol is more
2 suitable for transmission over a satellite link than using a TCP protocol.

1 5. The apparatus of claim 2 wherein said converting comprises
2 removing said header to leave said data substantially intact.

1 6. The apparatus of claim 2 wherein said converting comprises
2 removing said header to leave said data substantially intact and encapsulating said data
3 using a satellite protocol header.

1 7. The apparatus of claim 6 wherein said data is a portion of said flow
2 of information.

1 8. The apparatus of claim 1 wherein said processor is further
2 operatively disposed to receive said flow of information by a gateway over said
3 telecommunications link.

1 9. A communication apparatus comprising:
2 a TCP interface;
3 a satellite gateway interface;
4 a system memory;
5 a bus interconnecting said TCP interface, said satellite gateway interface
6 and said system memory with a processor, said processor operatively disposed to:
7 intercept a first communication connection between a client and a server;
8 form a second communication connection between a first satellite gateway
9 and a second satellite gateway that is over a satellite link;
10 transmit information describing said first connection to said second
11 satellite gateway; and
12 form a third communication connection between said second satellite
13 gateway and a destination server using said information describing said first connection
14 wherein said forming said second connection and forming said third connection occur
15 transparently to said client and said server.

1 10. The apparatus of claim 9 wherein said information comprises a
2 client address and a destination server address.

1 11. The apparatus of claim 9 wherein said processor is further
2 operatively disposed to transmit a response from said second satellite gateway to said first
3 satellite gateway when said third communication connection with said destination server
4 occurs.

1 12. The apparatus of claim 9 wherein said processor is further
2 operatively disposed to transmit a response from said first satellite gateway to said client
3 when said third communication connection with said destination server occurs.

1 13. The apparatus of claim 9 wherein said processor is further
2 operatively disposed to transmit a failure response from said first satellite gateway to said
3 client when said third communication connection is lost.

1 14. A apparatus for establishing a communication between a first
2 computer and one of a plurality of second computers, said apparatus comprising:
3 a network interface;
4 a satellite gateway interface;
5 a system memory;
6 a bus interconnecting said network interface, said satellite gateway
7 interface and said system memory with a processor, said processor operatively disposed
8 to:

9 provide a first protocol and a second protocol, wherein at least one of said
10 plurality of second computers is able to communicate with said first computer using said
11 first protocol and at least one of said plurality of second computers is able to
12 communicate with said first computer using said second protocol;

13 determine whether a connection can be established between said first
14 computer and at least one of said plurality of second computers using said first protocol;
15 if said connection cannot be established, then establish a connection
16 between said first computer and said at least one of said plurality of second computers
17 using said second protocol.

1 15. The apparatus of claim 14 wherein each of said plurality of second
2 computers is able to communicate with said first computer using said second protocol.

1 16. The apparatus of claim 14 wherein said first protocol is XTP.

1 17. The apparatus of claim 14 wherein said second protocol is TCP/IP.

1 18. The apparatus of claim 14 wherein said first protocol has a first
2 throughput and said second protocol has a second throughput, said first throughput being
3 at least 7.5 times greater than said second throughput at a bit error rate of 1×10^{-7}

1 19. The apparatus of claim 14 wherein said first protocol has a first
2 throughput and said second protocol has a second throughput, said first throughput being
3 at least 10 times greater than said second throughput at a bit error rate of 1×10^{-6}

1 20. The apparatus of claim 14 wherein said first protocol has a
2 throughput of at least 95% of an available bandwidth at a bit error rate of 1×10^{-8} .

1 21. In a communications system for providing transport of packetized
2 information in a TCP protocol between a client and a server, a method of converting said
3 system to a system for providing transport of packetized information in a satellite
4 protocol, said method comprising:

5 installing a first gateway, said first gateway operatively disposed to
6 intercept connections from said client to said server; said first gateway further operatively
7 disposed to convert said packetized information from TCP protocol to satellite protocol;

8 installing a second gateway, said second gateway disposed to establish a
9 connection with said server, said second gateway further operatively disposed to convert
10 said packetized information from satellite protocol to TCP protocol, said first gateway
11 and said second gateway operatively disposed to establish a connection over a common
12 telecommunications link.

INTERNET OVER SATELLITE APPARATUS**ABSTRACT OF THE DISCLOSURE**

5 According to the present invention a telecommunications apparatus for providing transport of packetized information over large distances. The apparatus provides a bi-directional flow of information from a first satellite gateway using a connection over a satellite network to a second satellite gateway. The apparatus uses a gateway that translates the information using the TCP protocol into information using a
10 satellite protocol, which is suitable for transmission of such information over the satellite network.





FIG. 1



F16. 2




FIG. 2A




FIG. 3A




FIG. 3B




Fig. 3C




Fig. 3D




Fig. 3E






Fig. 4



Delay vs. Throughput

FIG. 5



Bit Error Rate vs. Throughput (10Mbps Link)

FIG. 6




Fig. 7

DECLARATION AND POWER OF ATTORNEY

As a below named inventor, I declare that:

My residence, post office address and citizenship are as stated below next to my name; I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural inventors are named below) of the subject matter which is claimed and for which a patent is sought on the invention entitled: **INTERNET OVER SATELLITE APPARATUS** the specification of which X is attached hereto or _____ was filed on _____ as Application No. _____ and was amended on _____ (if applicable).

I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above. I acknowledge the duty to disclose information which is material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56. I claim foreign priority benefits under Title 35, United States Code, Section 119 of any foreign application(s) for patent or inventor's certificate listed below and have also identified below any foreign application for patent or inventor's certificate having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Country	Application No.	Date of Filing	Priority Claimed Under 35 USC 119

I hereby claim the benefit under Title 35, United States Code § 119(e) of any United States provisional application(s) listed below:

Application No.	Filing Date
60/118,227	February 2, 1999

I claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code, Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56 which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

Application No.	Date of Filing	Status

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith.

Richard T. Ogawa, Reg. No. 37,692

Kenneth R. Allen, Reg. No. 27,301

Lesley S. Craig, Reg. No. 26,578

Darin J. Gibby, Reg. No. 38,464

Chad S. Hilyard, Reg. No. 40,647

Roger T. Barrett, Reg. No. 41,599

William F. Vobach, Reg. No. 39,411

Scott L. Ausenblus, Reg. No. 42,271

Thomas D. Franklin, Reg. No. 43,616

Send Correspondence to:

Roger T. Barrett

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834

Direct Telephone Calls to:
 (Name, Reg. No., Telephone No.)




Name: **Roger T. Barrett**

Reg. No.: **41,599**

Telephone: **303-571-4000**

Full Name of Inventor 1:	Last Name: TOPOREK	First Name: JEROME	Middle Name or Initial: D.
Residence & Citizenship:	City: Venice	State/Foreign Country: California	Country of Citizenship: United States
Post Office Address:	Post Office Address: 343 Rennie Avenue, #8	City: Venice	State/Country: California Postal Code: 90291
Full Name of Inventor 2:	Last Name: PALTER	First Name: DAVID	Middle Name or Initial: C.
Residence & Citizenship:	City: Los Angeles	State/Foreign Country: California	Country of Citizenship: United States
Post Office Address:	Post Office Address: 10933 Wellworth Avenue, #3	City: Los Angeles	State/Country: California Postal Code: 90024
Full Name of Inventor 3:	Last Name: MCCOOEY	First Name: JEREMY	Middle Name or Initial: A.
Residence & Citizenship:	City: Los Angeles	State/Foreign Country: California	Country of Citizenship: United States
Post Office Address:	Post Office Address: 1002 Timberton Avenue, #215 820 Hilgard Avenue, #215 (In)	City: Los Angeles	State/Country: California Postal Code: 90024
Full Name of Inventor 4:	Last Name: HASSON	First Name: MARC	Middle Name or Initial: B.
Residence & Citizenship:	City: Redondo Beach	State/Foreign Country: California	Country of Citizenship: United States
Post Office Address:	Post Office Address: 1722 Goodman Avenue	City: Redondo Beach	State/Country: California Postal Code: 90278
Full Name of Inventor 5:	Last Name: HARTRICK	First Name: TIMOTHY	Middle Name or Initial: W.
Residence & Citizenship:	City: Los Angeles	State/Foreign Country: California	Country of Citizenship: United States
Post Office Address:	Post Office Address: 11965 Dorothy 10969 Rochester Avenue, #205 (In)	City: Los Angeles	State/Country: California Postal Code: 90047
Full Name of Inventor 6:	Last Name: GUYER	First Name: KAY	Middle Name or Initial: A.
Residence & Citizenship:	City: Los Angeles	State/Foreign Country: California	Country of Citizenship: United States
Post Office Address:	Post Office Address: 10933 Wellworth Avenue, #11	City: Los Angeles	State/Country: California Postal Code: 90024

I further declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

Signature of Inventor 1  JEROME D. TOPOREK Date <u>26 Jan 00</u>	Signature of Inventor 2  DAVID C. PALTER Date <u>Jan 26 2000</u>	Signature of Inventor 3  JEREMY A. MCCOOEY Date <u>January 26, 2000</u>
Signature of Inventor 4  MARC B. HASSON Date <u>1/26/00</u>	Signature of Inventor 5  TIMOTHY W. HARTRICK Date <u>1/26/00</u>	Signature of Inventor 6  KAYLA GUYER Date <u>January 26, 2000</u>